

Multi-Sensor Event Detection using Shape Histograms

Ehtesham Hassan
ehtesham.hassan@tcs.com

Gautam Shroff
gautam.shroff@tcs.com

Puneet Agarwal
puneet.a@tcs.com

TCS Research, New Delhi, India

August 19, 2014

Abstract

Vehicular sensor data consists of multiple time-series arising from a number of sensors. Using such multi-sensor data we would like to detect occurrences of specific *events* that vehicles encounter, e.g., corresponding to particular maneuvers that a vehicle makes or conditions that it encounters. Events are characterized by similar waveform patterns re-appearing within one or more sensors. Further such patterns can be of variable duration. In this work, we propose a method for detecting such events in time-series data using a novel feature descriptor motivated by similar ideas in image processing. We define the *shape histogram*: a constant dimension descriptor that nevertheless captures patterns of *variable* duration. We demonstrate the efficacy of using shape histograms as features to detect events in an SVM-based, multi-sensor, supervised learning scenario, i.e., multiple time-series are used to detect an event. We present results on real-life vehicular sensor data and show that our technique performs better than available pattern detection implementations on our data, and that it can also be used to combine features from multiple sensors resulting in better accuracy than using any single sensor. Since previous work on pattern detection in time-series has been in the single series context, we also present results using our technique on multiple standard time-series datasets and show that it is the most versatile in terms of how it ranks compared to other published results.

0.1 Introduction

While engineering new vehicles quality engineers schedule trial runs in a controlled environment to simulate the life cycle of a vehicle in the field. In order to do this effectively, it is important to understand how similar models are used in the field by customers, especially in terms of the distribution of certain maneuvers that stress various components, such as rapid acceleration, sudden braking, sharp turns, etc. The problem reduces to detecting occurrences of specific events from time-series data obtained from multiple on-board sensors across numerous runs of a large population of vehicles.

Training data was generated by running test vehicles in a controlled environment where the maneuvers of interest are performed and their time-stamps are logged. The task is to detect similar events in large collections of time-series data coming from vehicles in the field, efficiently and with high accuracy. In current practice engineers use hand-crafted rules to detect such events, for example “if sensor-1 is higher than t_1 and derivative of sensor-2 is more than t_2 ”, occurrence of an event is assumed to be indicated. However, it was found that such rules could not discover many occurrences of actual events: for example, we observed only 60-70% accuracy when using the rule based approach.

Pattern recognition in time series, even in the single sensor case, is a well-studied research problem. The problem at hand is further complicated because of the following practical challenges as observed in real-life data: a) the duration of an event is not fixed, b) the magnitude of the events may vary, c) sensor readings are often sparse, d) if the event is divided into sub-sections, the ratio of the durations of these sub-sections varies in different instances of the same event, and e) these events need to be detected from large time-series, therefore linear time algorithm is required. It is because of such challenges that the available techniques were found not to perform as well as needed, in terms of accuracy, efficiency or both.

The problem can be viewed as that of supervised classification in interesting and non-interesting categories since known occurrences of the specific patterns are available. However, previous work on pattern detection in time-series [7, 9, 26, 30] have primarily been on single time-series and also do not address some of the above challenges. In this paper we present a method for detection of time-series events that addresses these challenges.

We introduce a novel descriptor, the *shape histogram*, for representing of variable length time-series subsequences by exploiting the *shape* of a temporal pattern. The concept is borrowed from similar ideas in image processing, applied here in the one-dimensional world of time-series analysis. The shape histogram generates a time-scale invariant representation of time-series subsequences and is conveniently usable with various distance measures and learning algorithms. We have applied our proposed shape histogram to the specific task of event detection from time-series by using it as a feature vector for supervised classification.

Modeling time-series events requires exploiting long range correlations together with the short range interactions within subsequences. Our shape histogram descriptor addresses this problem by defining a shape-based representation of patterns that represents variable duration subsequences by a two-dimensional vector. The shape-histogram is a robust representation of time-series subsequences tolerant to variable time-scales as well as missing values.

Shape-histogram-based event detection is independent of expansions, compressions and discontinuities due to the inherent non-linearity embedded in the subsequence generation process in comparison with existing works which are applicable to constant duration time-series patterns. Also, the feature vector derived from a shape histogram can be easily combined with other features such as the duration or amplitude of an event. Thus we were able to use multiple sensors for our classification task, which results in improvements over any single sensor approach. We also demonstrate that our shape histogram based approach is most versatile and has the best rank as compared to various approaches available in research literature, on multiple standard time-series datasets.

The time complexity of generating a shape histogram from a time-series subsequence is quadratic in the size (duration) of the subsequence. However, subsequences, albeit of variable length, are always much smaller than the length of the time series and can be assumed to be bounded by a constant. Computing the shape histogram for all subsequences in a time-series is thus a linear-time procedure. Since the size of the shape-histogram is fixed and small, a discriminative classifier, such as SVM, also works efficiently.

Lastly, our method of feature extraction is un-supervised and performs better in both efficiency and quality even as compared to supervised methods of feature extraction such as shapelets[30, 20, 18] that exploit knowledge of class labels while deriving features, which we do not assume.

The document is structured as follows: We begin in Section 0.2 with an overview of our definition of multi-sensor time-series events and challenges faced in practice with vehicular sensor data. We motivate and introduce our novel shape-histogram feature descriptor in Section 0.3. Experimental evaluation of our technique on a real-life vehicle sensor data as well as standard datasets is presented in Section 0.4. Section 0.5 describes related work in time-series pattern detection.




Figure 1: An example time series generated with random values and two example subsequences

0.2 Overview

Previous results in time series mining can be broadly categorized in two groups. The first is to discover patterns (motifs) in an unsupervised manner, often using clustering of time-series subsequences[22]. The second set of approaches exploit domain knowledge and available training data for supervised classification of subsequences in order to detect specific events characterized by particular waveform patterns[18]. We are concerned with the latter problem.

Here we formally define a time-series, and a subsequence. A time series t_i is an ordered set of m values, i.e., $t_i = \{q_1, \dots, q_m\}$. The sequential order of values characterizes their temporal ordering with respect to a predefined sampling rule. Often there are multiple sensors in a system and more than one of the sensors may characterize the target event. Therefore we assume that there are multiple such time-series corresponding to different sensors, i.e., $T_i = \{\tau_i^1, \dots, \tau_i^j, \dots, \tau_i^s\}$ for s sensors. We get such a time-series from every session of operation of the underlying system, i.e., one set of time-series T_i for every independent run of a vehicle. Assuming there are n such runs / time-series, we have time-series dataset $T = \{T_1, \dots, T_i, \dots, T_N\}$. For single sensor time-series we use a shorthand τ_i^j , which refers to a time-series from i^{th} run of a vehicle and j^{th} sensor.

A subsequence S is a smaller section of τ_i^j . Formally, a subsequence S consists of a set of *approximately* l contiguous positions from τ_i^j , i.e., $S = \{q_i, q_{i+1}, \dots, q_{i+l-1}\}$. We also assume that length of such subsequences is not fixed, i.e., the length of subsequences $l = |S|$ can vary, albeit being bounded by a maximum value considerably smaller than the length of each time series. This definition captures the general case where no condition is imposed on the length of each subsequence or the overlap between two subsequences (refer Figure 1). In this setting we define the problem of event detection in time-series as follows:

Given a training set of (potentially variable length) subsequences labeled as interesting or non-interesting, i.e., $\{1, 0\}$, the objective is to classify a set of unlabeled candidate subsequences (also of variable length) that have been extracted from a set of time-series T .

0.2.1 Solution Overview

The first step of the solution is related to identification of candidate subsequences from the time-series dataset T . Most previous works have used constant window size subsequences, which requires some way to choose the correct or optimum window size within which all occurrences of the desired patterns will fall. However, in practice, the patterns characterizing an interesting event are present in subsequences of different length that are all nevertheless visually similar.

For defining the boundary of a subsequence we assume that hand-crafted rules based on domain-specific heuristics are available, which is indeed the case in practice. We illustrate the use of such rules in the case of time-series obtained from multiple sensors available in a commercial vehicle operating in varying operating conditions. Details of the dataset is described in Section 0.4, and our method of identifying candidate subsequences, and an example of vehicular event *Event1* is given in Appendix-A.




Figure 2: Potential occurrences of *Event1* in *Sensor2*.

Figure 2 depicts multiple occurrences of the same event, but of variable time duration. Further, in many practical




Figure 3: Sensor reading (y-axis) for *Event1*, time on x-axis.

scenarios, real-time measurements have missing values because sensors sometimes fail to capture data, especially for extreme values, as can be observed in Figure 3 taken from our vehicle sensor data. Such regions of missing values also occur multiple times and are of varying duration within an event.

After extraction of candidate subsequences we construct a feature vector for every subsequence using the shape-histograms (described in Section 0.3). A key property of the feature vector drawn from shape histogram is that it can be combined with shape histogram of other sensors. For this we collate subsequences from other sensors sharing same time-stamp and concatenate their shape-histogram features vectors, as well as add additional features if desired. Finally we train an SVM classifier using the training data and use it to classify the candidate subsequences into interesting and uninteresting categories.

0.3 Shape histogram

0.3.1 Motivation

Developments in image processing and computer vision have produced a variety of shape descriptors that exploit distinct attributes such as shape, appearance, and texture to develop feature-based object representation. Some of the well established descriptors are discussed in [3, 4, 19, 15, 28, 8] have been shown to result in robust object representations that display low intra-class and high inter-class variance in different application scenarios.




Figure 4: Time-series subsequence S plotted in \mathbf{R}^2

In time-series, however, patterns are uniquely characterized by the sequential order of numerical or symbolic attributes, where the underlying temporal order between consecutive values defines a highly correlated structure. Each point q_i of a time-series subsequence S is defined by two attributes: 1) its value (w), and 2) its temporal order (t), i.e., $q_i = (w_i, t_i)$. Figure 4 shows an example subsequence in which two points q_i and q_j are highlighted. The value w , and temporal order t of a point q_i can be treated as coordinates of a point in two-dimensional Euclidean space,

in which the subsequence defines a curve. The intuition is that curves of similar shape in this two-dimensional space should correspond to subsequences representing the same event. Motivated by the success of shape based object descriptions in image processing, such as the ‘shape context’¹[3], we define a similar representation for time-series patterns that we call the *shape histogram*.

0.3.2 Definition

We assume the prior availability of subsequences filtered from a large time-series via rule based techniques as described in Section 0.2.1 and Appendix-A.

The shape histogram (H) for a subsequence S captures the distribution of all pairs of points (i.e., in two dimensional value, time space) with respect to their relative distances and angles. For example, in Figure 4, the relative angle of q_j from q_i is θ , computed with respect to the X-axis with origin at q_i . Similarly the distance of q_j from q_i is the euclidian distance in two dimensional value-time space. We compute the angle and distance between each pair $(q_i, q_j), q_j \neq q_i$. *Note:* One may think that it is sufficient to use either w or t instead of using both in euclidean distance, because the other dimension is based on angular distance. However, after a careful observation of extrema cases ($w_i - w_j = 0$ OR $t_i - t_j = 0$) it can be understood that it is important to use both in the distance dimension.

These relative pairwise angles and distances encode the trend in sensor values with respect to time and form a two-dimensional distribution which is the shape histogram. Given a subsequence $S = \{q_1, \dots, q_i, \dots, q_l\}$ we define its multi-dimensional *shape histogram* $H \in \mathbb{N}^{m \times n}$ as follows:

$$H(\mathbf{k}) = [\#\{(q_i, q_j) \in \text{bin}(\mathbf{k})\} | (q_i \neq q_j)] \quad (1)$$

Here $\mathbf{k} \in \mathbb{N}^2$ and $H(\mathbf{k})$ is the \mathbf{k}^{th} values of the shape histogram H , and it is the number of pairs (q_i, q_j) that have the relative angular and euclidean distances within a given range, i.e., that fall in $\text{bin}(\mathbf{k})$, which defines the distance and angle ranges of \mathbf{k}^{th} bin of the histogram. We shall describe how the ranges for each bin are computed in Section 0.3.3.

Note that some shape descriptors, also based on shape context, compute the distances in log-polar space [3, 8]. However, we want the shape histogram for a time-series subsequence to be equally sensitive to near and distantly placed points, as the dominance of local versus global shape properties varies from pattern to pattern. Therefore, the proposed shape histogram is based directly on the distance and angle rather than log-polar coordinates. We also assume the points (q_i) are uniformly spaced on the time-scale.

0.3.3 Computation

The shape histogram in the context of time-series is a two dimensional histogram, with the dimensions being along the distance bins, and angle bins. We assume it has m distance bins and n angle bins. Sample shape histograms are shown in Figure 6.

For each point $q_i \in S$, we compute its distance and angle with the remaining points of S in two-dimensional Euclidean space. In the angle computation for q_i , origin is set at q_i , and orientation of remaining points is computed with respect to the X axis. Figure 4 shows point q_j located at the angle θ_{ij} from q_i , computed with origin at q_i . The distance between points $q_i \equiv (w_i, t_i)$ and $q_j \equiv (w_j, t_j)$ is defined as $d_{ij} = \sqrt{(w_i - w_j)^2 + (t_i - t_j)^2}$.

Repeating the procedure for all points, we get the distance matrix $D = [d_{ij}]$ and angle matrix $A = [\theta_{ij}]$ where each row corresponds to a point q_i in S . We divide each row of the distance matrix by maximum distance in that row to normalize it to a $[0, 1]$ scale. The angles are anyway defined on a fixed scale of 0 to 359.99 degrees so they do not require normalization. The shape histogram H for a subsequence S will be a two-dimensional histogram representing the joint distribution of elements in D and A , computed as follows.

A shape histogram calculated for all pairs of points, keeping one of the points q_i fixed, also has same size $(m \times n)$ and is referred as the *shape context* of the point q_i . First, we compute the shape context H_i of each point q_i defined by its distance-angle based distribution with respect to the remaining points. Considering $[d_0, d_1] \cup [d_1, d_2] \dots \cup [d_{m-1}, d_m]$ as the bins for distance quantization, and $[\alpha_0, \alpha_1] \cup [\alpha_1, \alpha_2] \dots \cup [\alpha_{n-1}, \alpha_n]$ as the bins for angle quantization, i.e., quantizing the distances in m levels and angles in n levels. (Note that all distance bins, and angle bins are of uniform size.) For point q_i , the count of the $\text{bin}(x, y)$ of its shape context is defined as

$$H_i(x, y) = \sum_{j=1, j \neq i}^l \phi(D_{i,j}, A_{i,j}, d_{x-1}, d_x, \alpha_{y-1}, \alpha_y), \quad (2)$$

¹http://en.wikipedia.org/wiki/Shape_context

where

$$\begin{aligned} & \phi(D_{ij}, A_{i,j}, d_{x-1}, d_x, \alpha_{y-1}, \alpha_y) \\ &= \begin{cases} 1 & \text{if } D_{i,j} \in [d_{x-1}, d_x], A_{i,j} \in [\alpha_{y-1}, \alpha_y] \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

Summing the shape contexts of all points in S i.e., $H_i(:, :) \forall i \in 1, \dots, l$ defines the global distribution of the point-pairs. The summation $H_{sum} = \sum_i H_i$ represents distribution of euclidean distances and angles between the points in S . In H_{sum} , count of each $bin(x, y)$ represents the count of pair points, which are arranged within distance $[d_{x-1}, d_x]$ and angle $[\alpha_{y-1}, \alpha_y]$. Dividing H_{sum} by its maximum entry we arrive at the shape histogram H for subsequence S .

Next, the value of every bin of H is then identified as a feature for discriminative classification, therefore we convert the shape histogram H into $p = m \times n$ dimensional feature vector.

Further, a very important property of the shape histograms is that they can be combined with other features extracted from the same time-series, such as average value \bar{w}_i of the subsequence. Not just this, additional features can be drawn from the shape histograms of other sensors of the same system, and can be combined in order to improve the accuracy of time-series event detection. This property facilitates the usage of a discriminative classifier and we call this method of combining the shape histograms with other features or other shape histograms as *linear concatenation*. We demonstrate this aspect of shape histograms through experiments on a real-life dataset in Section 0.4.2.




Figure 5: Contrived example of subsequences ($S_1 - S_6$) and surface plots of corresponding shape histogram ($H_1 - H_6$) placed below. Each subsequence is of different length limited in the underneath brace. (Axis values can be ignored.)

0.3.4 Examples and Intuition

Notice that H is invariant to $l = (|S|)$ making it applicable for time-series subsequences of variable length. Figure 5 shows shape histogram for some example subsequences to establish their scale invariance. In other words, the point-pair distribution histogram encodes the temporal arrangement of variables in a subsequence, and provides easy access to inherent semantic information embedded in the shape, characterized by the temporal order. Similar shapes generate point-pair distribution histograms having similar slope characteristics. Also, it must be observed that shape histogram is invariant to small shifts in onset of interesting state transitions. Small values of m , and n have been used in Figure 5 for explanatory purposes, in practice higher values are used that give better results.

The pattern of time-series points exhibit strong correlation. Though small amount of stochastic element is present which makes future values only partially related to past values. Nevertheless, for the subsequences characterizing a particular event or motif, the relative arrangement of histogram points is unique.




Figure 6: Two subsequences described in Figure 1 and corresponding shape histogram (Subsequence 1 and Subsequence 2 shown in Figure 1) and corresponding shape histogram

The global shape of a subsequence described by distantly positioned point-pairs is identified by the overall spread of the surface of H . The local shape of a subsequence is described by peaks in H , as can be observed in Figure 5.

Parameters (m, n) control the robustness and discriminative ability of the shape histogram. In the supervised learning setting, parameter search by splitting the complete training data in training and validation set is the standard approach. However, this method is computationally complex, and requires sufficient labeled examples. Instead, we have evaluated the shape histogram for range of distance and orientation quantization levels, and the best results are presented. Initial selection of (m, n) is based on the heuristic that a lower number of bins do not capture the uniqueness of point-pair distribution. Therefore, we select sufficiently large number of angular and distance bins such that the descriptor accurately captures the discriminative pattern of the point distribution histograms. The experimental evaluation showed that with increase in shape histogram parameters, the discriminative ability of feature increases with increased similarity matching cost. However, for very large number of bins (m, n) , sparsely distributed point distribution histograms are noisy and highly sensitive intra class variations. This is illustrated in Figure 6 depicting a shape histogram using more realistic values of parameters.

0.3.5 Time and space complexity

In a subsequence S having l values, we have lC_2 unique-pairs requiring equal number of distance and angle computations for shape histogram computation which are easy to compute. Shape histogram is of the dimension $m \times n$ requiring proportional storage requirement, which is typically a small number. Therefore, while our method of using shape histogram has time complexity quadratic in the length of the candidate subsequences, these are themselves of bounded length as compared to the length of the series itself. Thus our approach is linear in the length of the overall time series.

Further, the shape histogram is a low-dimensional feature vector that can be classified efficiently by standard classifiers. Additionally, the histogram often has significant number of empty bins by the example shown in Figure 6 and substantiated by experimental results in the next Section. This property also presents opportunities for further optimization from the efficiency perspective.

0.4 Experimental Results

To substantiate the key contributions of our work, we present the results from three different kinds of experiments, performed on two types of datasets. First one was performed on a real-life dataset, obtained from the sensors of vehicles. This dataset results in variable length subsequences, with intermittently missing values. We present

the results of this analysis in Section 0.4.1. Later, in Section 0.4.2, we present a summary of results indicating that it is possible to concatenate shape histogram of other sensors of the same electro-mechanical system, and that accuracy of event detection can be improved through this. Thereafter, in Section 0.4.3, we present a summary and analysis of results obtained from the experiments performed on 24 different public datasets. Through these last set of experiments we demonstrate that our approach of using shape-histograms and classification mechanism such as SVM with RBF kernel, gives better results than many other approaches published in the research literature.

0.4.1 Evaluation on vehicular sensor data

We used our proposed shape histogram technique to detect events of interest from our vehicular sensor data as introduced in Section 0.1. Our data collection consisted of measurements from seventeen different sensors. *Event1* is characterized by dynamic behaviour of different sensors, nevertheless, we begin with a single sensor results using *Sensor2* (speed) so as to illustrate the behaviour of our technique. Subsequently, we present results using other sensors as well, as well as results using multiple sensors together, which significantly improves performance.

Table 1: *Event1* identification using *Sensor2*

shape histogram		
Accuracy	Precision	Recall
96.52	94.41	88.05
shapelet ²		
89.3	81.04	85.18

First 831 candidate subsequences potentially containing occurrences of *Event1* were extracted from each sensor using the rule-based approach as described in Appendix-A. These rules were relaxed in order to get 100% recall, and the rules were drawn based on domain knowledge. Here, by relaxed rules we mean, adjusting the threshold values in the rule, e.g., if one of the rules that characterizes the event is stated as, speed should be more than about 30 miles per hour, we could use 20 miles per hour to ensure 100% recall. The 831 subsequences varied in length from 9 to 39 time-steps. From this candidate set, 122 subsequences were manually tagged as containing occurrences of *Event1* and 709 subsequences as false candidates, i.e., these did not contain a genuine occurrence of *Event1*.

We first present results of event detection for *Event1* using a single sensor, *Sensor2*. This and all further experiments have been performed via 10-fold cross-validation using the manually tagged data as described above. A range of histogram parameters were tried out; in practice 15 distance bins and 40 orientation bins achieved the best results. Initial evaluation with a nearest neighbour classifier using the euclidean distance metric between shape histograms, or even a linear SVM classifier did not improve over results obtained via the rule based approach. However, using a radial (RBF) kernel with the SVM classifier showed significant improvement in the classification accuracy. We therefore, use Shape-Histograms + SVM with RBF kernel (ShapeHist-SVM-RBF) to benchmark the results of our approach with other approaches.

In order to compare our approach with other approaches, we have also applied shapelets². [30] for classification of subsequences, for the same experimental setting. Table 1 shows results of both the methods. As is evident, our approach of shape histograms coupled with SVM classifier using RBF kernel(ShapeHist-SVM-RBF) performs significantly better, at least on this particular time-series and event. We analyze the reasons for this improvement below.

Figure 7 shows examples of *Event1* occurrences missed by shapelet and correctly identified by shape histogram. In general, we observed that shapelet failed in identifying the events generated by extreme braking conditions having low true positives and high true negatives (positives are the genuine *Event1* events). An actual *Event1* is generally characterized by high *Sensor1*, and fast decrease in *Sensor2* occurring for small time-duration. However, exceptional situations of prolonged engine acceleration may increase the duration of *Event1* as shown in Figure 7.

The shapelet based approach concentrates on a small local region, i.e., the region depicting the transition (shown by a red ellipse) for classification of subsequences. However, large intra-class variation does not provide sufficient distinction between categories in the decision tree architecture used by shapelet. Also, the decision about the occurrence of an event also requires the information of initial and final state of the sensor variable. In these cases, our shape histogram exploits the combination of local information, i.e., transition stage and global information i.e., initial and final state.

We also observed some cases where the shape histogram approach wrongly classified the events, most often they were characterized by relatively smooth decrease in *Sensor2* measurement, i.e., a false *Event1*. In general, we desire high recall in such problems because of subtle difference in patterns of real and false *Event1*.

². The shapelet code used for our experiments was taken from <http://www.cs.ucr.edu/~mueen/LogicalShapelet/>




Figure 7: Wrongly classified examples of *Event1*, by Shapelets, on subsequences of *Sensor2*, shown by a red oval

Table 2: *Event1* detection: multiple sensors

Sensor	Accuracy	Precision	Recall	Dist. bins	Ori. bins
<i>Sensor3</i>	92.78	92.10	92.80	15	90
<i>Sensor4</i>	95.05	94.80	95.10	25	20
<i>Sensor5</i>	91.31	92.10	91.30	30	25
<i>Sensor6</i>	92.25	91.40	92.20	30	20
<i>Multi-Sensor</i>	97.19	97.10	97.20	-	-

0.4.2 Linear Concatenation of Features

As mentioned earlier and illustrated in Appendix-A, occurrences of *Event1* actually involve the dynamics of multiple sensors. This requirement in real-life dataset and the fact that shape histograms of multiple sensors can be combined, make the use of shape histogram for detection of events in vehicular sensor data even more appropriate. Therefore, we also evaluated the shape histogram approach for identifying *Event1* from other sensors including *Sensor4*, *Sensor5*, and *Sensor6*. These sensors represent acceleration, brake pressure, and deceleration measurements in the vehicle dataset. The results presented in Table 2 used experimentally obtained optimal shape histogram parameters for each of these sensors as shown in Table 2.

Different sensors respond for a specific event with different dynamics, nevertheless, complete characterization of the event requires information of complex interaction between responses of different sensor. Therefore we subsequently combined the shape histograms for all the sensors by linear concatenation. As also shown in Table 2, identification of *Event1* by combining multiple shape histograms for all these sensors showed 0.67% gain in the identification accuracy. These results demonstrate that shape histograms can exploit information from multiple sensors to improve classification accuracy.

0.4.3 Evaluation on public datasets

For further evaluation our ShapeHist-SVM-RBF approach, we collected multiple datasets available in public domain. We took all of the available³ datasets from [18], and 7 additional datasets from [14]. We present accuracy, precision, and recall for our approach on these datasets in Table 4 (Appendix-B). Number of Distance bins (m), Orientation bins (n), and average occupancy has also been reported in the same table.

It may appear that with the above parameter settings that the shape histogram for each subsequence is a vector having 1500 (25×60) real values. However, in practice on average only 4.08% of total bins (Occupancy) across the instances have non-zeros entities which is of the order actual subsequence length; as a result storing and computing distances between shape histogram feature descriptors can be done efficiently.

To arrive at the results presented in Table 4, we experimented with a range of quantization levels for distances and orientations. First, a euclidean-distance-based nearest neighbor classifier as well as a linear kernel based support vector machine (SVM) was applied. In this case, shape histograms with 25 distance bins and 60 angle bins achieved better results in comparison with other parameter settings. However, using the L2-norm based similarity metric and SVM based classification using a linear decision boundary were less effective. Similar to how it worked on the

³ Out of many datasets used in [18], only 17 were available.

Table 3: Testing Accuracy taken from [18], ShapeHist-SVM-RBF included with new ranks.

Dataset Name	ShapeHist		SVM Linear		Rotation Forest		Random Forest		Bayesian Net-work		Naive Bayes		1-NN		C4.5		Shapelet Tree	
Adiac	28.25	4.0	23.79	9.0	30.69	1.0	30.43	2.0	25.06	7.0	28.13	5.0	25.32	6.0	24.3	8.0	29.92	3.0
Beef	73.33	4.5	86.67	2.0	70.00	6.0	60.00	7.5	90.00	1.0	73.33	4.5	83.33	3.0	60.00	7.5	50.00	9.0
Chlorine-Concen.	64.18	1.0	56.15	8.0	63.52	2.0	57.58	4.0	57.08	5.0	45.96	9.0	56.93	6.0	56.48	7.0	58.8	3.0
Coffee	100	2.5	100	2.5	89.29	8	100	2.5	96.43	5.5	92.86	7.0	100	2.5	85.71	9.0	96.43	5.5
DiatomSize-Red.	91.88	3.0	92.16	2.0	83.01	5.0	80.39	6.0	90.20	4.0	78.76	7.0	93.46	1.0	75.16	8.0	72.22	9.0
ECGFiveDays	99.60	1.0	98.95	3.0	98.61	4.0	93.26	8.0	99.54	2.0	96.4	6.0	98.37	5.0	96.17	7.0	77.47	9.0
Face (four)	96.73	6.0	97.73	4.5	98.86	3.0	87.5	7.0	100	1.5	97.73	4.5	100	1.5	76.14	9.0	84.09	8.0
GunPoint	98.67	3.5	100	1.0	98.67	3.5	96.00	6.0	99.33	2.0	92.00	7.0	98.00	5.0	90.67	8.0	89.33	9.0
ItalyPower-Demand	94.28	1.0	92.13	5.5	92.03	7	93.00	2.0	92.42	4.0	92.52	3.0	92.13	5.5	90.96	8.0	89.21	9.0
Lighting7	64.38	4.5	69.86	1.0	65.75	2.5	64.38	4.5	65.75	2.5	57.53	6.0	49.32	8.5	53.42	7.0	49.32	8.5
Medical-Images	58.98	1.0	52.5	2.0	51.45	3.0	50.79	4.0	28.16	8.0	17.37	9.0	45.66	6.0	44.87	7.0	48.82	5.0
Mote-Strain	82.23	9.0	88.66	4.0	86.98	5.0	84.58	6.0	89.06	2.0	88.82	3.0	90.34	1.0	84.42	7.0	82.51	8.0
Sony-AIBORob	91.54	1.0	86.69	4.0	89.02	3.0	85.19	5.0	89.68	2.0	79.03	9.0	84.03	8.0	84.53	6.5	84.53	6.5
Symbols	64.97	8.0	84.62	4.0	84.42	4.0	84.62	4.0	92.26	1.0	77.99	6.5	85.63	2.0	47.14	9.0	77.99	6.5
Synthetic-Cont.	77.33	8.0	87.33	6.0	92.00	3.0	89.00	5.0	76.67	9.0	78.00	7.0	93.00	2.0	90.33	4.0	94.33	1.0
Trace	100	1.5	98.00	6.0	98.00	6.0	98.00	6.0	100	1.5	98.00	6.0	98.00	6.0	98.00	6.0	98.00	6.0
TwoLead-ECG	99.56	1.0	99.30	3.0	97.98	6.0	96.14	7.0	98.77	5.0	99.12	4.0	99.47	2.0	85.25	8.0	85.07	9.0
Average Rank	3.56		3.97		4.24		5.09		3.71		6.09		4.18		7.41		6.76	

real-life dataset, the best classification results were observed using SVM with the Gaussian kernel with standard deviation = 3 and cost parameter $C = 3.33$. The results corresponding these parameters are shown in Table 4.

Further, to compare the results of our approach with other approaches presented in the research literature, we took the table (Table 3 of [18]) that reports the accuracy of various approaches in [18] on multiple datasets. This is because, to the best of our knowledge work done by Lines et al. in [18] by and large reports the best accuracies of various approaches of event detection on time-series. To execute this comparison we took only the 17 datasets³ that were common between our and their experiments, included our results (ShapeHist-SVM-RBF) in the table, and recalculated the ranks as shown in Table 3. For calculation of these ranks we assigned ‘average ranks’ to the algorithms that perform equally well, e.g., in Coffee dataset, 4 approaches give 100% accuracy, we assign them all a rank of $(1 + 2 + 3 + 4)/4$. Unlike [18], our method of assigning these ranks is same as how ranks are assigned in Wilcoxon signed rank test[27]. We also performed the Wilcoxon signed rank test and were able to reject the null hypothesis for Naive-Bayes, Shapelets, and C4.5. However, when we performed z-test on the two most competing approaches (ShapeHist and Bayesian-Network), we were able to reject the null hypothesis indicating that the mean of the two approaches are not same.

We plotted a comparison of our approach with the average accuracy of various approaches, on the 17 common³ datasets as shown in Figure 8(Appendix-B). It can be observed from this plot that our approach performs consistently better than other approaches barring a few cases. Similarly, we also plotted a comparison of our approach with the best accuracy of any approach reported in [18], on the 17 common³ datasets also shown in Figure 9. From this, we observe that our approach is almost similar to the best performing approach on multiple datasets. (Note: Single approach that performs the best on all these datasets does not exist.) Also, on other datasets such as SwedishLeaf and Lighting2 that are not included in [18], our approach performs better. Further, from these observations and from Table 3, it can be inferred that by and large **the average rank of our approach is 3.56, which is the best, making it the most versatile approach for detection of events from time-series data.**

0.5 Related work

Detecting events of interest in time-series data has been studied in many different application domains. Initial work explored a variety of methods and classifiers based on statistics and machine learning theory. Guralnik and Srivastava [7] addressed the time series event detection as a change detection problem by modeling the time-series using a piecewise segmented model. However, they focus on change point detection and divide the whole time-series in different segments based on these change points. Our problem is different in that our events of interest are more complex than change points alone, e.g., they may involve steady regions as well as sharp dips, etc.

Early work by Rodriguez et al. [24] introduced inductive logic programming for multivariate time-series classification as well as the application of boosting for performance improvement of the classifiers. Their approach

worked with variable length event patterns however it required expert intervention for definition of a region. A knowledge-based event detection framework for time-series arising in health monitoring systems is discussed in [9]. Other approaches include using recurrent neural networks for time-series classification by organizing the network dynamics in alignment with class labels[10]. Availability of sufficient labeled data is important aspect of time-series classification problem. In this context, Wei and Keogh [26] presented semi-supervised learning using nearest neighbor time-series classification with insufficient labeled data. However, all these approaches work with fixed length subsequences only.

Xi et al. [29] proposed a dynamic-time-warping-based nearest neighbor classifier for addressing time complexity in case of large datasets. Recently Jeong et al. [12] presented weighted dynamic time warping for time-series classification by associating higher weights to phase differences in local neighborhood. Such approach suffer from issues of efficiency and it would be hard to process large time-series datasets. In summary, each of the above methods concentrated on experimentation with different classifiers and similarity measures. However, they did not adequately address real-time requirements of time and space complexity.

In [30], Ye and Keogh introduced ‘shapelet’ based time-series classification, where a shapelet is a time-series sub-sequence considered as maximal representative of a class. Information-gain based shapelet discovery locates the most informative local pattern for classification using linear search over the complete time-series. Subsequently a decision tree is defined on the computed information gains to perform the time-series classification. This approach suffers from high time-complexity.

In an extension of shapelet based time-series analysis, Mueen et al. [20] proposed logical combination of multiple shapelets for performance improvement. Lines et al. [18] extended the shapelet concept for transforming the time-series in uniform data space so as to be able to use other classifiers in addition to decision trees. The k best shapelets extracted from the data in combination with decision tree classifier were shown to achieve comparable performance with the results presented in [30]. In further extensions of the shapelet idea of local pattern analysis, Zakaria et al. [31] proposed unsupervised learning based shapelet discovery method for clustering time-series subsequences. The underlying philosophy is the creation of min-dist based distance map between a local pattern (read unsupervised shapelet), and a set of time-series subsequences.

Most of the above pattern detection techniques deal with detecting fixed-length waveform patterns. Shapelets, however, do treat variable-length patterns as does our shape histogram approach; therefore our experimental comparisons in Section 0.4 are primarily with the shapelet approach, especially on our real-life vehicular data where variable length events are indeed present. At the same time, since the shapelet descriptors are themselves of variable length this restricts their use in conjunction with state-of-the art machine learning classifiers such as SVMs; in contrast, shape histograms do not have this restriction. (Note that this advantage that is also shared by the shapelet-transform extension by Lines et al. [18].) Also, the shapelet discovery algorithms discussed in [30, 20, 18] require apriori label information: in other words, they require labelled data to calculate the shapelet descriptors themselves, rather than just for the classification stage. In contrast our shape histogram feature descriptor can be computed independent of the events to be detected, i.e., un-supervised feature extraction, and still performs better on real data.

0.6 Conclusions

We have proposed the shape histogram, a novel feature descriptor for waveform patterns inspired by ideas from image processing and computer vision, and shown that it is useful for event identification in time-series. The shape histogram is able to tackle scenarios involving varying duration events as well as is able to handle sparsity in data caused by missing values. We have presented experimental results on real-life vehicular sensor data and shown that our approach is superior to available implementations of previous pattern detection approaches. We have also evaluated our technique using publicly available data and shown that it achieves results comparable with the best previous results on those datasets. Our method presents novel extension of shape based feature extraction techniques to time-series data representation and demonstrates the applicability of shape-based features derived from image processing for event detection in real-life time-series data, especially that arising in the context of vehicular sensor analysis.

Bibliography

- [1] R. J. Alcock, Y. Manolopoulos, D. E. Laboratory, and D. O. Informatics. Time-series similarity queries employing a feature-based approach. In *In 7 th Hellenic Conference on Informatics, Ioannina*, pages 27–29, 1999.
- [2] A. Bagnall, L. M. Davis, J. Hills, and J. Lines. Transformation based ensembles for time series classification. In *SDM*, pages 307–318, 2012.
- [3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 24(24):509–522, April 2002.
- [4] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05) - Volume 1*, CVPR ’05, pages 886–893, 2005.
- [5] D. Eads, D. Hill, S. Davis, S. Perkins, J. Ma, R. Porter, and J. Theiler. Genetic Algorithms and Support Vector Machines for Time Series Classification. In *Proc. SPIE 4787; Fifth Conference on the Applications and Science of Neural Networks, Fuzzy Systems, and Evolutionary Computation; Signal Processing Section; Annual Meeting of SPIE*. 2002.
- [6] S. Gudmundsson, T. Runarsson, and S. Sigurdsson. Support vector machines and dynamic time warping for time series. In *Neural Networks, 2008. IJCNN 2008. (IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on*, pages 2772–2776, June 2008.
- [7] V. Guralnik and J. Srivastava. Event detection from time series data. In *Proceedings of the 5th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 33–42, 1999.
- [8] E. Hassan, S. Chaudhury, and M. Gopal. Word shape descriptor-based document image indexing: a new dbh-based approach. *International Journal on Document Analysis and Recognition (IJDAR)*, 16(3):227–246, 2013.
- [9] J. Hunter and N. Mcintosh. Knowledge-based event detection in complex time series data. In *Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making*, pages 271–280. Springer-Verlag, 1999.
- [10] M. Husken and P. Stagge. Recurrent neural networks for time series classification. *Neurocomputing*, 50:223–235, 2003.
- [11] A. Jalba, M. Wilkinson, J. Roerdink, M. Bayer, and S. Juggins. Automatic diatom identification using contour analysis by morphological curvature scale spaces. *Machine Vision and Applications*, 16(4):217–228, 2005.
- [12] Y.-S. Jeong, M. K. Jeong, and O. A. Omitaomu. Weighted dynamic time warping for time series classification. *Pattern Recognition*, 44(9):2231 – 2240, 2011.
- [13] E. Keogh and A. Ratanamahatana. Everything you know about dynamic time warping is wrong. *3rd Workshop on Mining Temporal and Sequential Data, in conjunction with 10th ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD-2004)*, Seattle, WA, 2004.
- [14] E. J. Keogh, Q. Zhu, B. Hu, X. Hao, Y. and Xi, L. Wei, and C. A. Ratanamahatana. The ucr time series classification/clustering homepage:www.cs.ucr.edu/~eamonn/time_series_data/. 2011.
- [15] I. Laptev. On space-time interest points. *International Journal of Computer Vision*, 64:107–123, September 2005.

- [16] L. Li, J. McCann, N. S. Pollard, and C. Faloutsos. Dynammo: Mining and summarization of coevolving sequences with missing values. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '09, pages 507–516, 2009.
- [17] J. Lin, E. Keogh, L. Wei, and S. Lonardi. Experiencing sax: A novel symbolic representation of time series. *Data Min. Knowl. Discov.*, 15(2):107–144, 2007.
- [18] J. Lines, L. M. Davis, J. Hills, and A. Bagnall. A shapelet transform for time series classification. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 289–297, 2012.
- [19] D. G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision*, volume 2, pages 1150–1157, 1999.
- [20] A. Mueen, E. Keogh, and N. Young. Logical-shapelets: an expressive primitive for time series classification. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1154–1162, 2011.
- [21] R. T. Olszewski. *Generalized Feature Extraction for Structural Pattern Recognition in Time-series Data*. PhD thesis, Pittsburgh, PA, USA, 2001. AAI3040489.
- [22] T. Rakthanmanon, E. J. Keogh, S. Lonardi, and S. Evans. Time series epenthesis: Clustering time series streams requires ignoring some data. In *Proceedings of the 2011 IEEE 11th International Conference on Data Mining*, pages 547–556, 2011.
- [23] C. A. Ratanamahatana and E. J. Keogh. Making time-series classification more accurate using learned constraints. In M. W. Berry, U. Dayal, C. Kamath, and D. B. Skillicorn, editors, *SDM*. SIAM, 2004.
- [24] J. J. Rodríguez, C. A. González, and H. Boström. Learning First Order Logic Time Series Classifiers: Rules and Boosting. In *PKDD '00: Proceedings of the 4th European Conference on Principles of Data Mining and Knowledge Discovery*, pages 299–308. Springer-Verlag, 2000.
- [25] D. Roverso. Multivariate temporal classification by windowed wavelet decomposition and recurrent neural networks. In *In 3rd ANS International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface*, 2000.
- [26] L. Wei and E. Keogh. Semi-supervised time series classification. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '06, pages 748–753, 2006.
- [27] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics*, 1(6):80–83, 1945.
- [28] G. Willems, T. Tuytelaars, and L. Gool. An efficient dense and scale-invariant spatio-temporal interest point detector. In *Proceedings of the 10th European Conference on Computer Vision: Part II*, pages 650–663, 2008.
- [29] X. Xi, E. Keogh, C. Shelton, L. Wei, and C. A. Ratanamahatana. Fast time series classification using numerosity reduction. In *Proceedings of the 23rd international conference on Machine learning*, pages 1033–1040, 2006.
- [30] L. Ye and E. Keogh. Time series shapelets: a new primitive for data mining. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 947–956, 2009.
- [31] J. Zakaria, A. Mueen, and E. Keogh. Clustering time series using unsupervised-shapelets. 2012.

Appendix - A

An interesting event in vehicular sensor data might be as described colloquially by an engineer as follows: “*Event 1* is characterized by a sudden high deceleration of a vehicle running at high speed reaching an almost stationary position within 2-3 seconds after applying breaks.” A hand-crafted rule for detecting the point of exact state transition, i.e., critical point during *Event1* might characterize such an event as follows:

$$\{(Sensor1 \text{ is greater than } s_1) \cap (Sensor2 \text{ more than } s_2) \cap (Sensor3 \text{ is in } [s_{31}, \dots, s_{3m}]) \cap (\text{drop in } Sensor2 \text{ is more than } s_2^*) \}$$

Sensor1, *Sensor2*, and *Sensor3* refer values of brake cylinder pressure, wheel speed and gear as sampled in real-time by onboard sensors. Subsequently, heuristic rules help us identify the left and right boundary of these critical points, to identify subsequences that contain occurrence of *Event1*. For example, we might define the extreme boundaries of the event sequence in terms of the following conditions:

Left boundary: $\{(Sensor1 \text{ is less than } a \text{ for consecutive } b \text{ seconds}) \cup (c \text{ seconds before the critical point})\}$

Right boundary: $\{(Sensor1 \text{ is less than } d \text{ for consequent } e \text{ seconds}) \cap (Sensor2 \text{ in more than } f \text{ for consecutive } g \text{ seconds}) \cup (35 \text{ Points before the critical point})\}$

Figure 2 depicts time series data of *Sensor2* for two runs of a vehicle along with occurrences of possible *Event1* defined in rectangular windows. When such an approach based on hand-crafted rules were used for event detection, we found that it gave poor results; e.g., we observed barely 60-70 % accuracy for *Event1*. Further, different sets of rules are used by each engineer, which makes it difficult to arrive at principled conclusions based on analysis of field data.

However, if instead we use rules similar to those above but with more relaxed conditions (i.e., different, less conservative values of $s_1, s_2, s_3, a, b, c, d, e, f, g$), we get can 100% recall, but very poor precision: Exactly such relaxed rules are used to generate *candidate subsequences*, as referred to in Sections 2.1 and 3.2, on which our shape-histogram-based event detection procedure is then applied.

Appendix - B




Figure 8: Average accuracy of various approaches reported in [18], compared with the accuracy of ShapeHist-SVM-RBF; on 17 common³ datasets.




Figure 9: The best accuracy of various approaches reported in [18], compared with the accuracy of ShapeHist-SVM-RBF; on 17 common³ datasets.

Table 4: Time-series classification Shape histogram on standard datasets

S.N.	Dataset	Distance bins (m)	Orientation bins (n)	Accuracy	Precision	Recall	Average Occupancy
1.	Adiac [11]	25	36	28.25	28.44	28.18	5.63
2.	Beef [2]	45	10	73.33	74.64	73.68	13.10
3.	ChlorineConcentration [16]	30	30	64.18	60.64	63.97	9.12
4.	Coffee [18]	35	45	100	100	100	7.27
5.	DiatomSizeReduction [14]	32	40	91.88	92.34	91.12	6.79
6.	ECG [21, 30, 18]	40	45	91.35	90.64	90.87	3.76
7.	ECGFiveDays [21, 30, 18]	27	40	99.60	99.37	99.48	6.14
8.	Face (four) [13]	36	90	96.73	95.83	97.12	5.16
9.	FacesUCR [14]	25	60	86.56	86.44	85.66	4.14
10.	GunPoint [30, 23]	25	60	98.67	97.63	97.63	4.06
11.	ItalyPowerDemand [18]	36	36	94.28	94.26	94.18	4.86
12.	Lighting2 [17, 6, 5]	40	36	88.64	88.91	87.64	8.61
13.	Lighting7 [17, 6, 5]	40	40	64.38	65.11	64.76	4.03
14.	MedicalImages [14]	35	36	58.98	56.46	58.10	6.38
15.	MoteStrain [30, 18]	25	45	82.23	82.73	82.08	5.37
16.	OliveOil [2]	40	36	80.00	79.84	80.13	10.48
17.	SonyAIBORobot Surface [14]	35	15	91.54	89.91	91.46	8.23
18.	SonyAIBORobot SurfaceII [14]	25	10	88.88	89.11	89.02	9.66
19.	SwedishLeaf [29]	20	60	74.53	74.78	74.24	6.41
20.	SyntheticControl [1]	30	60	77.33	78.10	78.63	3.88
21.	Symbols [14]	35	36	64.97	64.76	64.39	7.94
22.	Trace [23, 25, 14]	35	36	100	100	100	7.64
23.	TwoLeadECG [21, 30, 18]	36	36	99.56	99.45	99.58	4.45
24.	Wafer [26, 14]	50	45	99.68	99.90	99.58	6.87